

PATENT
IBM Docket No. GB9-2000-0032

Listing of Claims (indicating status and amendments):

1 1. (Original) A method for recovering from failures affecting a resource manager within a
2 group of resource managers, wherein the resource managers within the group have access to a
3 shared resource via which remote resource managers communicate with the resource managers
4 within the group, the shared resource including data storage structures to which resource
5 managers within said group connect to send and receive communications, the method
6 comprising:

7 storing, within a first data storage structure of the shared resource, unit of work
8 descriptors for operations performed in relation to said shared resource by the resource managers
9 in said group;

10 sending a notification of a connection failure between a second data storage structure of
11 the shared resource and a first resource manager within said group, the notification being sent to
12 the remaining resource managers within the group which are connected to the second data
13 storage structure;

14 one or more of said remaining resource managers accessing said first data storage
15 structure and analysing the unit of work descriptors to identify the units of work relating to the
16 second data storage structure that were being performed by the first resource manager when the
17 connection failure occurred; and

18 said one or more remaining resource managers recovering the identified units of work.

1 2. (Original) A method according to claim 1 wherein, if there are no remaining resource
2 managers connected to the second data storage structure after said connection failure, said

PATENT
IBM Docket No. GB9-2000-0032

3 notification is sent to a remaining resource manager when that resource manager connects to the
4 second data storage structure.

1 3.(Original) A method according to claim 1 wherein, if there are no remaining resource
2 managers connected to the second data storage structure after said connection failure, the failed
3 resource manager determines when it is restarted whether any other resource manager has
4 performed recovery for its units of work relating to the second data storage structure and, upon
5 determining that no resource manager has performed said recovery, the restarted resource
6 manager recovers said units of work.

1 4. (Original) A method according to claim 1, wherein all remaining resource managers within
2 the group which are connected to the second data storage structure respond to said notification by
3 attempting to access said first data storage structure to identify units of work to recover, and the
4 method includes the further steps of:

5 responsive to a first remaining resource manager identifying a unit of work to recover,
6 said first remaining resource manager attempting to set a flag for said unit of work;

7 responsive to successfully setting said flag, assigning recovery responsibility for said unit
8 of work to said first remaining resource manager; and

9 refusing to assign recovery responsibility for said unit of work to said first remaining
10 resource manager if said flag has been set by another remaining resource manager.

1 5. (Original) A method according to claim 4, including the further step of:

PATENT
IBM Docket No. GB9-2000-0032

2 responsive to said flag having been set by another remaining resource manager, said first
3 remaining resource manager attempting to identify a further unit of work to recover and
4 attempting to set a flag for said identified further unit of work.

1 6.(Original) A method according to claim 4, including the following steps in response to a
2 connection failure between the second data storage structure of the shared resource and said first
3 remaining resource manager during recovery of said unit of work:

4 sending a notification of said connection failure to the remaining resource managers
5 within the group which are connected to the second data storage structure;

6 one or more of said remaining resource managers accessing said first data storage
7 structure and analysing the unit of work descriptors to identify the units of work relating to the
8 second data storage structure that were being performed by the first remaining resource manager
9 when the connection failure occurred; and

10 said one or more remaining resource managers recovering the identified units of work

1 7. (Original) A method according to claim 1, wherein the unit of work descriptors include:

2 a unit of work identifier;

3 an identification of messages put or retrieved within the unit of work;

4 a status for the unit of work; and

5 a sequence number.

PATENT
IBM Docket No. GB9-2000-0032

8. (Original) A method according to claim 1, wherein the shared resource is a coupling facility list structure, the second data storage structure is a coupling facility list structure in which a coupling facility list header represents a shared access message queue, and the first data storage structure is an administration list structure of the coupling facility for storing unit of work descriptors.

9. (Original) A method according to claim 8, including storing within the coupling facility, for each resource manager within the group, a list header information map representing the set of shared access message queues within the second data storage structure for which the resource manager has performed some work.

10. (Original) A method according to claim 9, including reading said list header information map during recovery to identify the set of shared access message queues within the second data storage structure for which the failed resource manager has performed some work.

11. (Original) A method according to claim 1, including storing within the shared resource a structure interest map identifying the set of data storage structures to which respective resource managers within said group are connected.

12. (Original) A method according to claim 11, wherein the step of recovering the identified units of work is a first recovery phase and wherein the method includes a second recovery phase comprising the steps of:

reading the structure interest map for the failed resource manager to identify the set of data storage structures to which the failed resource manager was connected at the time of said connection failure;

PATENT
IBM Docket No. GB9-2000-0032

A 7 identifying any operations performed by the failed resource manager on said set of data
8 storage structures which were not recovered in the first recovery phase; and
9 one or more of said remaining resource managers then backing out said unrecovered
10 operations.

1 13. (Original) A method according to claim 12, wherein the method includes setting a key for
2 operations performed in relation to the shared resource, the key identifying the resource manager
3 which performed the operation, and wherein the identification of operations performed by the
4 failed resource manager comprises checking said keys for unrecovered operations performed in
5 relation to any of said set of data storage structures.

1 14. (Original) A method according to claim 1, wherein a single unit of work represented by a
2 unit of work descriptor may include operations performed in relation to a plurality of data storage
3 structures, and wherein the partial units of work corresponding to said operations are recovered
4 by different ones of said remaining resource managers within the group.

1 15. (Original) A method for recovering from failures affecting a resource manager within a
2 group of resource managers, wherein the resource managers within the group have access to a
3 shared resource, the shared resource including data storage structures to which resource managers
4 within said group connect to perform operations in relation to data held in said shared resource,
5 the method comprising:

6 storing, within a first data storage structure of the shared resource, unit of work
7 descriptors for operations performed by the resource managers in said group in relation to data
8 held in said shared resource;

PATENT
IBM Docket No. GB9-2000-0032

9 sending a notification of a connection failure between a second data storage structure of
10 the shared resource and a first resource manager within said group, the notification being sent to
11 the remaining resource managers within the group which are connected to the second data
12 storage structure;

13 one or more of said remaining resource managers accessing said first data storage
14 structure and analysing the unit of work descriptors to identify the units of work relating to the
15 second data storage structure that were being performed by the first resource manager when the
16 connection failure occurred; and

17 said one or more remaining resource managers recovering the identified units of work.

1 16. (Original) A method according to claim 15, wherein the data storage structures of said
2 shared resource include data storage structures which contain shared message queues and said
3 operations performed in relation to said shared resource include putting messages onto a shared
4 message queue and retrieving messages from a shared message queue, for communication
5 between a remote resource manager and resource managers within said group.

1 17. (Original) A method according to claim 16, wherein the unit of work descriptors include:

2 a unit of work identifier;

3 an identification of messages put or retrieved within the unit of work;

4 a status for the unit of work; and

5 a sequence number.

PATENT
IBM Docket No. GB9-2000-0032

A/ 1 18. (Original) A method according to claim 16, wherein the operations of putting messages onto
2 a shared queue and retrieving messages from a shared queue are performed under transactional
3 scope such that a message which is put is only available to resource managers other than the
4 resource manager putting the message after commitment of the put operation and a message
5 which is retrieved is only available to the retrieving resource manager after commitment of the
6 retrieval operation, and wherein said stored unit of work descriptors identify each of the
7 following:

8 units of work that were uncommitted but for which a decision to commit had been made
9 when the failure occurred;

10 units of work that were uncommitted but for which a decision to abort had been made
11 when the failure occurred; and

12 units of work for which no commit or abort decision had been made when the failure
13 occurred;
14 and wherein recovering the identified units of work comprises:

15 committing message put and retrieval operations for which a decision to commit had been
16 made;

17 backing out message put and retrieval operations for which a decision to back out had
18 been made; and

19 backing out message put and message retrieval operations for which no commit or abort
20 decision had been made.

PATENT
IBM Docket No. GB9-2000-0032

19. (Original) A distributed data processing system including:

a plurality of resource managers;

a shared access resource including data storage structures to which the resource managers connect to send and receive communications to and from remote resource managers, the shared access resource including:

means for storing, within a first data storage structure of the shared resource, unit of work descriptors for operations performed in relation to said shared resource by the resource managers in said plurality; and

means for sending a notification of a connection failure between a second data storage structure of the shared resource and a first resource manager within said plurality, the notification being sent to the remaining resource managers within the plurality which are connected to the second data storage structure; wherein said remaining resource managers include:

means for accessing said first data storage structure and analysing the unit of work descriptors to identify the units of work relating to the second data storage structure that were being performed by the first resource manager when the connection failure occurred; and

means for recovering the identified units of work.

20. (Original) A computer program product comprising program code recorded on a machine-readable recording medium, the program code comprising the following set of components:

PATENT
IBM Docket No. GB9-2000-0032

A1
3 a plurality of resource managers;

4 a shared access resource manager including program code for managing storage and
5 retrieval of data within data storage structures to which the resource managers connect to send
6 and receive communications to and from remote resource managers, the shared access resource
7 manager including:

8 means for storing, within a first data storage structure of the shared resource, unit of work
9 descriptors for operations performed in relation to said shared resource by the resource managers
10 in said plurality; and

11 means for sending a notification of a connection failure between a second data storage
12 structure of the shared resource and a first resource manager within said plurality, the notification
13 being sent to the remaining resource managers within the plurality which are connected to the
14 second data storage structure;

15 wherein said remaining resource managers include:

16 means for accessing said first data storage structure and analysing the unit of work
17 descriptors to identify the units of work relating to the second data storage structure that were
18 being performed by the first resource manager when the connection failure occurred; and

19 means for recovering the identified units of work.